arXiv:1403.5491v1 [math.PR] 21 Mar 2014

On trees invariant under edge contraction

Olivier Hénard* Pascal Maillard†

March 24, 2014

Abstract

We study random trees which are invariant in law under the operation of contracting each edge independently with probability $p \in (0,1)$. We show that all such trees can be constructed through Poissonian sampling from random measured real trees satisfying a natural scale invariance property. This has connections to exchangeable partially ordered sets, real-valued self-similar increasing processes and quasi-stationary distributions of Galton–Watson processes.

1 Introduction

Take a random rooted tree T (in the graph sense) and contract each edge independently with probability $p \in (0,1)$. Are there (necessarily infinite) random trees which are invariant in law under this operation? Trivial examples are the (semi)-infinite ray, i.e. the one-dimensional half-lattice $\mathbb{N}$ and with root 0, or a (possibly random) number of copies of this graph whose roots are identified. Slightly less trivial examples are the previous graph with a bouquet of edges attached to each vertex, the number of edges in each bouquet being iid according to a geometric distribution (starting at 0). In this article, we construct/characterize all possible examples in terms of certain *(measured) real trees* arising as scaling limits of these trees; the trees then can be obtained from their limit by a certain Poissonian sampling procedure. We also consider the special cases of trees which are invariant (in law) under translations along the infinite ray, this has connections with real-valued self-similar increasing processes and with quasi-stationary distributions of subcritical Galton–Watson processes.

The study of this problem originates in a geometrically motivated question asked to us by I. Benjamini: Is it possible to find a law on the space of trees and a suitable renormalization procedure, i.e. a suitable random coarse-graining operation that preserves the tree structure, such that the law of this tree is invariant under this operation? The contraction operation is an example of such a renormalization procedure, which has the pleasant feature of allowing to characterize all random trees (with finitely many ends) it leaves invariant.

We are not aware of any similar results in the random tree literature. The operations that are usually considered on trees, for example random growth or pruning, always operate on leaves or whole subtrees instead of single interior vertices, see [Rém, AP98]. There might be a good reason for that: The trees we obtain are indeed very different from usual trees in the sense that they are very elongated, with long chains of vertices of degree 2, to which might be attached some bouquets of edges as in the above example. In particular, neither exponentially growing trees nor critical Galton–Watson trees (conditioned on non-extinction) are amongst them.

The current article however has strong connections to exchangeable random partially ordered sets studied by Janson [Jan11] (see Remark 1.13). Indeed, the ancestral relation of a tree defines a partial order on the set of vertices; contracting edges then amounts to restricting the partial

*School of Mathematical Sciences, Queen Mary University of London, Mile End Road, London E1 4NS, U.K.
O.H. is a postdoctoral research assistant supported by EPSRC grant EP/J004022/2
†Faculty of Mathematics, Weizmann Institute of Science, POB 26, Rehovot 76100, Israel. Partially supported by a grant from the Israel Science Foundation

order to a subset of the vertices. The ancestral relations of the limiting real trees then can be represented as exchangeable partial orders on $\mathbb{N}$, which yields powerful tools for studying them. The current study therefore falls into the general program initiated by Aldous [Ald93] of using exchangeability to study limits of random discrete structures; for examples of other settings see the references in [Jan11] or [HP11].

We now define the space of discrete trees that we will be working with. We consider rooted, locally finite trees $T = (V, E, \rho)$ in the graph-theoretic sense; in combinatorics these are also known as *unlabeled*, *unorderered*, *non-plane* or *Pólya* trees [Sta97, FS09, Drm09]. We further require that $T = (V, E, \rho)$ has finitely many distinct ends (defined as non backtracking infinite paths from the root). We say that two such trees are *equivalent* if there exists a root-preserving graph isomorphism between them. We then denote by $\mathbb{T}$ the space of equivalence classes of trees and by $\mathbb{T}_f \subset \mathbb{T}$ the subspace of finite trees. Note that for simplicity, we will always identify an equivalence class with its representatives. We endow the space $\mathbb{T}_f$ with the discrete topology and the space $\mathbb{T}$ with the topology of local convergence defined as follows: For a tree $T \in \mathbb{T}$ and $k \in \mathbb{N}$, denote by $T^{\leqslant k}$ its restriction to the vertices at (graph) distance at most k from the root. A sequence $(T_n)_{n\geqslant 0}$ in $\mathbb{T}$ then is said to converge to $T \in \mathbb{T}$ if $(T_n^{\leqslant k})_{n\geqslant 0}$ converges to $T^{\leqslant k}$ in $\mathbb{T}_f$. It is easy to see that this topology is metrizable in such a way that the space $\mathbb{T}$ is a complete separable metric space. Note also that the induced topology on $\mathbb{T}_f$ is indeed the discrete topology.

A tree $T = (V, E, \rho) \in \mathbb{T}$ determines a partial order $\preceq_T$ on its vertex set V by $v \preceq w$ if and only if v lies on the shortest path from ρ to w. We also write $v \precnsim_T w$ if and only if $v \preceq_T w$ and $v \neq w$. If $v \precnsim_T w$, we also say that v is an *ancestor* of w, or that w is a *descendant* of v; accordingly, we call $\preceq_T$ the *ancestral relation*. Note that the tree T can in fact be completely recovered from its ancestral relation. This allows us to formally define the contraction operation as follows:

Definition 1.1. Let $T = (V, E, \rho) \in \mathbb{T}$ and let $V' \subset V$ be a subset of its vertices containing the root. We then define the *contracted tree* $\mathcal{C}(T, V')$ as follows: This is the rooted tree T' with vertex set V', root ρ and whose ancestral relation $\preceq_{T'}$ is the restriction of $\preceq_T$ to the vertex set V':

$$\mathcal{C}(T, V') = T'$$

In order to deal with *random* trees, we endow the space $\mathbb{T}$ with the Borel-σ-field induced by its topology. We then denote by $\mathcal{M}_1(\mathbb{T})$ the space of probability measures on $\mathbb{T}$, endowed with the topology of weak convergence. Note that there exists a canonical embedding of $\mathbb{T}$ into $\mathcal{M}_1(\mathbb{T})$ given by $T \mapsto \delta_T$. For this reason, we will often denote the elements of $\mathcal{M}_1(\mathbb{T})$ by T as well, and will refer to them as *random* trees. We similarly define $\mathcal{M}_1(\mathbb{T}_f)$.

We now define a certain contraction operation from $\mathbb{T}$ to $\mathcal{M}_1(\mathbb{T})$. For a tree $T \in \mathbb{T}$, denote by $\mathrm{Spine}(T)$ the subset of its vertices (called the *spine*) which lie on an infinite path starting at ρ (including the root).

Definition 1.2. Let $p, q \in (0, 1]$. Let $T = (V, E, \rho) \in \mathbb{T}$. Set $V_0 = V \backslash \mathrm{Spine}(T)$ and $V_1 = \mathrm{Spine}(T) \backslash \rho$. The random tree $\mathcal{C}_{p,q}(T) \in \mathcal{M}_1(\mathbb{T})$ is defined to be equal to $\mathcal{C}(T, V')$, where V' is the random subset of vertices containing ρ, every vertex $v \in V_0$ independently with probability p and every vertex $w \in V_1$ independently with probability q. Formally, define a map $\varphi_T : \{0, 1\}^{V_0 \cup V_1} \to \mathbb{T}$ by

$$\varphi_T((x_v)_{v\in V_0\cup V_1}) = \mathcal{C}(T, \{\rho\} \cup \{v \in V_0 \cup V_1 : x_v = 1\}).$$

Then $\mathcal{C}_{p,q}(T)$ is defined to be the push-forward of the measure $(p\delta_1 + (1-p)\delta_0)^{\otimes V_0} \otimes (q\delta_1 + (1-q)\delta_0)^{\otimes V_1}$ by the map φ_T.

The map φ_T appearing in Definition 1.2 is measurable and the map $\mathcal{C}_{p,q} : \mathbb{T} \to \mathcal{M}_1(\mathbb{T})$ therefore is well defined. Furthermore, it is continuous and thus naturally extends to a continuous map from $\mathcal{M}_1(\mathbb{T})$ to itself by the mapping Theorem [Bil09, p.21]. We can now define *self-similar trees*.

Definition 1.3. Let $p, q \in (0, 1]$. We say that a random tree $T \in \mathcal{M}_1(\mathbb{T})$ is (p, q)-self-similar, if $T = \mathcal{C}_{p,q}(T)$.

We now extend the above definitions to $\mathbb{R}$-trees. A precise definition of $\mathbb{R}$-trees and of the Gromov–Hausdorff–Prokhorov (GHP) topology can be found in the appendix (Section A). We recall here that an $\mathbb{R}$-tree is a geodesic metric space $(\mathcal{V}, d)$ without "cycles", in the sense that for every $x, y \in \mathcal{V}$, if r and q are continuous injective maps from $[0,1]$ to $\mathcal{V}$ such that $q(0) = x$ and $q(1) = y$, and $r(0) = x$ and $r(1) = y$, then $q([0,1]) = r([0,1])$. A rooted, measured $\mathbb{R}$-tree is a quadruple $\mathcal{T} = (\mathcal{V}, d, \rho, \mu)$, where $\rho \in \mathcal{V}$ is a distinguished vertex called the *root*, and μ is a measure on $\mathcal{V}$. We say that two rooted, measured $\mathbb{R}$-trees are *equivalent* if there exists a root- and measure-preserving isometry between them. We then denote by $\mathfrak{T}$ the space of equivalence classes of locally compact rooted, measured $\mathbb{R}$-trees, such that the measure μ is boundedly finite, and the tree $\mathcal{T}$ has finitely many distinct ends (again, an end is defined as an infinite non backtracking path starting from the root). As above, we identify a tree with its equivalence class. We endow the space $\mathfrak{T}$ with the GHP-topology (see Section A) and with its induced Borel σ-field. We also denote by $\mathfrak{T}_f$ and $\mathfrak{T}_1$ the subspaces of trees, for which the measure μ is finite or a probability measure, respectively.

Every $\mathbb{R}$-tree $\mathcal{T}$ defines a *length measure* $\ell_{\mathcal{T}}$ on its set of vertices. In many cases of interest in probability theory, the length measure is not locally finite[1], for example in the case of Aldous's Brownian continuum random tree. However, in this article, we will be interested in those trees $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}$ for which the measure μ dominates the length measure $\ell_{\mathcal{T}}$, such that, in particular, the length measure is boundedly finite. We therefore define the spaces

$$\mathfrak{T}^\ell = \{\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T} : \mu \geqslant \ell_{\mathcal{T}}\}, \quad \mathfrak{T}_f^\ell = \mathfrak{T}_f \cap \mathfrak{T}^\ell, \quad \mathfrak{T}_1^\ell = \mathfrak{T}_1 \cap \mathfrak{T}^\ell.$$

We now define a discretization operation on the space $\mathfrak{T}^\ell$. For this, we recall that as for discrete trees, every rooted $\mathbb{R}$-tree $\mathcal{T} = (\mathcal{V}, d, \rho)$ induces a partial order $\preceq_{\mathcal{T}}$ on its set of vertices $\mathcal{V}$ by $x \preceq_{\mathcal{T}} y$ if and only if $x \in [\![\rho, y]\!]$, where $[\![\rho, y]\!]$ is the range of the geodesic from ρ to y (however, it is not true anymore that $\mathcal{T}$ can be recovered from $\preceq_T$). We can now define the following discretization operation:

Definition 1.4. Let $\mathcal{T} = (\mathcal{V}, d, \rho)$ be a rooted $\mathbb{R}$-tree. Let V_0 be a subset of $\mathcal{V}$ containing ρ and V_1 be a multiset of elements of $\mathcal{V}$, formally, this can be defined as a counting measure on $\mathcal{V}$. Suppose that V_0 and V_1 are boundedly finite in the sense that they contain a finite number of elements in each ball of finite radius. We then define the *discretized tree* $\mathcal{D}(T, V_0, V_1)$ as follows: We first define the rooted discrete tree T' with vertex set $V_0 \cup V_1$ and whose ancestral relation $\preceq_{T'}$ is defined as follows:

$$\forall v, w \in V_0 \cup V_1 : v \preceq_{T'} w \iff v \preceq_{\mathcal{T}} w \text{ and } v \in V_0.$$

We then set

$$\mathcal{D}(T, V_0, V_1) = T'$$

See Figure 1 for a graphical illustration.

Example 1.5. Consider the rooted $\mathbb{R}$-tree with only one element, the root. Define V_1 as the multiset which contains n times the root, and choose V_0 to be empty, $V_0 = \varnothing$. The vertices of the discretized tree $\mathcal{D}(T, V_0, V_1)$ are the root ρ plus n copies of it that we denote by $(\rho_i)_{1 \leqslant i \leqslant n}$, and the edges are the pairs $(\{\rho, \rho_i\},\ 1 \leqslant i \leqslant n)$.

We now define a randomized discretization operation $\mathcal{D} : \mathfrak{T}^\ell \to \mathcal{M}_1(\mathbb{T})$, which will turn out to play a fundamental role later on.

Definition 1.6. For a tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}^\ell$, define a random, rooted, discrete tree $\mathcal{D}(\mathcal{T}) \in \mathcal{M}_1(\mathbb{T})$ by $\mathcal{D}(\mathcal{T}) = \mathcal{D}(\mathcal{T}, V_0 \cup \{\rho\}, V_1)$, where

- V_0 is the set of atoms of a Poisson process on $\mathcal{V}$ with intensity measure $\ell_{\mathcal{T}}$, and
- V_1 is the multiset of atoms of a Poisson process on $\mathcal{V}$ with intensity measure $\mu - \ell_{\mathcal{T}}$.

[1] With respect to the topology on $\mathcal{V}$ induced by the metric d. One can define a finer topology generated by open segments with respect to which the length measure is always locally finite. Note that both topologies induce the same σ-algebra on the space of locally compact trees.

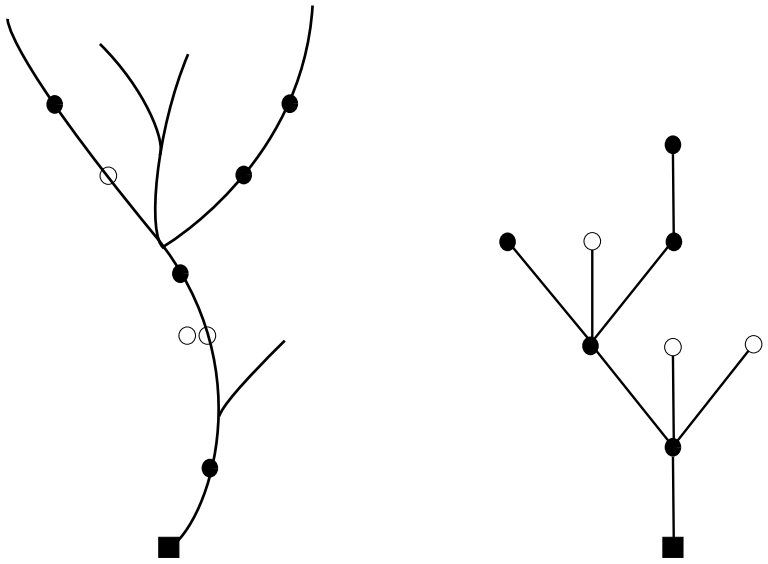

Figure 1: A real tree $\mathcal{T}$ (left) and a realization of its discretization $\mathcal{D}(\mathcal{T})$ (right). Black circles represent vertices sampled according to the length measure $\ell_{\mathcal{T}}$, while white circles represent vertices which have been sampled according to the measure $\mu - \ell_{\mathcal{T}}$. The root is represented by a square. The two neighboring white circles on the left indicate that two points have been sampled at the same spot, this is possible if the measure μ has atoms. As is visible from the right, white vertices are never an ancestor of other vertices.

The map $\mathcal{D} : \mathfrak{T}^\ell \to \mathcal{M}_1(\mathbb{T})$ from Definition 1.6 is well defined. The map $\mathcal{D}$ gives rise to a topology on the space $\mathfrak{T}^\ell$:

Definition 1.7. Define the topology $\mathcal{O}$ on $\mathfrak{T}^\ell$ to be the pull-back topology by the map $\mathcal{D}$, i.e. the coarsest topology on $\mathfrak{T}^\ell$ that makes the map $\mathcal{D}$ continuous.

We have the following result:

Theorem 1.8. *The topology $\mathcal{O}$ coincides with the GHP topology on $\mathfrak{T}^\ell$. In particular, the map $\mathcal{D}$ is continuous with respect to the GHP topology and is a stochastic kernel from $\mathfrak{T}^\ell$ to $\mathbb{T}$.*

Since the map $\mathcal{D}$ is a stochastic kernel from $\mathfrak{T}^\ell$ to $\mathbb{T}$ by Theorem 1.8, it extends to a map from $\mathcal{M}_1(\mathfrak{T}^\ell)$ to $\mathcal{M}_1(\mathbb{T})$, denoted by $\mathcal{D}$ as well.

As part of the proof of Theorem 1.8, we will establish the following lemma:

Lemma 1.9. *$\mathfrak{T}^\ell$, $\mathfrak{T}_f^\ell$ and $\mathfrak{T}_1^\ell$ are closed subspaces of $\mathfrak{T}$, $\mathfrak{T}_f$ and $\mathfrak{T}_1$, respectively. Furthermore, the space $\mathfrak{T}_1^\ell$ is compact.*

If the preceding definitions have seemed rather arbitrary and artificial, the next theorem will explain their meaning. But first, we have to define the notion of a *compatible sequence* of random trees.

Definition 1.10. For a random rooted tree $T = (V, E, \rho) \in \mathcal{M}_1(\mathbb{T}_f)$ with $|V| = n + 1$ almost surely, define for $m \leqslant n$, $\mathcal{C}(T, m) = \mathcal{C}(T, V' \cup \{\rho\})$, where V' is a uniformly chosen subset of $V \backslash \{\rho\}$ with m (distinct) elements. A family of random rooted trees $(T_n)_{n\in\mathbb{N}} \in \mathcal{M}_1(\mathbb{T}_f)$ with $|V(T_n)| = n + 1$ almost surely is called *compatible*, if for each $n \geqslant m \geqslant 1$, we have $T_m = \mathcal{C}(T_n, m)$.

Definition 1.11. For a random rooted measured $\mathbb{R}$-tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_f^\ell$, we define $\mathcal{D}(\mathcal{T}, n)$ to be the (random) tree $\mathcal{D}(\mathcal{T})$ conditioned on having $n + 1$ vertices.

Theorem 1.12. *There exists a one-to-one correspondence between compatible families of random rooted trees $(T_n)_{n\in\mathbb{N}}$ and random rooted measured $\mathbb{R}$-trees $\mathcal{T} \in \mathcal{M}_1(\mathfrak{T}_1^\ell)$, given by $T_n = \mathcal{D}(\mathcal{T}, n)$ for $n \geqslant 1$.*

Remark 1.13. The proof of the preceding theorem, as well as the proof of Theorem 1.8, relies very much on a representation of the compatible family $(T_n)_{n\in\mathbb{N}}$ and of the limiting tree $\mathcal{T}$ in terms of an exchangeable random partial order on $\mathbb{N}$. These have been studied in generality by Janson [Jan11], who provided a limiting representation of their *extremal* points in terms of a deterministic partial order on a continuous space. The contribution of our Theorem 1.12 for our setting is two-fold. First, we identify the limit as a tree from the space $\mathfrak{T}_1^\ell$, together with the right discretization/projection. Second, by the virtue of Theorem 1.8, we show that the construction of this limit is indeed measurable with respect to the Borel σ-field induced by the GHP topology on $\mathfrak{T}_1^\ell$, thus allowing us to speak of a *random* limit.

We now turn to self-similar $\mathbb{R}$-trees. For this, we define a shrinking operation $\mathcal{S}_{p,q}$ which will play the role of $\mathcal{C}_{p,q}$ for $\mathbb{R}$-trees. As in the case of discrete trees, for a tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}$ we denote by $\mathrm{Spine}(\mathcal{T})$ the subset of its vertices (called the *spine*) which lie on an infinite path starting at ρ (including the root). If $x \in \mathcal{V}$, we denote by $\vec{x}$ the most recent ancestor of x on the spine, i.e. the vertex in $[\![\rho, x]\!] \cap \mathrm{Spine}(\mathcal{T})$ with maximal distance from the root. For two vertices $x, y \in \mathcal{V}$, we then have

$$d(x, y) = d(x, \vec{x}) + d(\vec{x}, \vec{y}) + d(\vec{y}, y), \qquad \text{if } \vec{x} \neq \vec{y}.$$

Definition 1.14. If $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}^\ell$ and $p, q \in (0, 1]$, then we define the tree $\mathcal{T}' = \mathcal{S}_{p,q}(\mathcal{T}) = (\mathcal{V}', d', \rho', \mu')$ by

- $\mathcal{V}' = \mathcal{V}$ and $\rho' = \rho$,
- $d'(x, y) = pd(x, y) + (q - p)d(\vec{x}, \vec{y})$ and
- $\mu' = p\mu + (q - p)\ell_{\mathcal{T}}(\cdot \cap \mathrm{Spine}(\mathcal{T}))$.

In words, we shrink distances *off* the spine by a factor p and *on* the spine by a factor q.

The map $\mathcal{S}_{p,q} : \mathfrak{T}^\ell \to \mathfrak{T}^\ell$ is continuous and therefore naturally extends to a continuous map from $\mathcal{M}_1(\mathfrak{T}^\ell)$ to itself.

Definition 1.15. Let $p, q \in (0, 1]$. We say that a random rooted measured $\mathbb{R}$-tree $\mathcal{T} \in \mathcal{M}_1(\mathfrak{T}^\ell)$ is (p, q)-self-similar, if $\mathcal{T} = \mathcal{S}_{p,q}(\mathcal{T})$.

The operators $\mathcal{S}_{p,q}$, $\mathcal{C}_{p,q}$ and $\mathcal{D}$ are linked by the following commutation relation which follows from standard properties of Poisson processes.

Lemma 1.16. *We have the following equality of maps from $\mathfrak{T}^\ell$ to $\mathbb{T}$:*

$$\mathcal{D} \circ \mathcal{S}_{p,q} = \mathcal{C}_{p,q} \circ \mathcal{D}.$$

We can now finally state the main theorem of this article:

Theorem 1.17. *Let $p, q \in (0, 1]$. There exists a one-to-one correspondence between (p, q)-self-similar random rooted discrete trees $T \in \mathcal{M}_1(\mathbb{T})$ and (p, q)-self-similar random rooted measured $\mathbb{R}$-trees $\mathcal{T} \in \mathcal{M}_1(\mathfrak{T}^\ell)$ given by*

$$T = \mathcal{D}(\mathcal{T}).$$

This Theorem calls for a further study of the class of (p, q)-self-similar random rooted measured $\mathbb{R}$-trees. This study is initiated in Section 5, where we give an overview of the generality of examples that can be constructed. Very much like in the case of self-similar real-valued processe, it seems out of reach to completely characterize this family see [Ver85], at least without making further assumptions. Section 6 then discusses such assumptions. The first one is translation invariance:

if the (p,q)-self-similar random rooted discrete tree $T \in \mathcal{M}_1(\mathbb{T})$ is translation invariant, then so is the (p,q)-self-similar random rooted measured $\mathbb{R}$-tree according to Proposition 6.1. In this case, Proposition 6.3 asserts that only trivial $\mathbb{R}$-trees arise if $p = q$: this explains why we consider distinct contraction parameters p and q off the spine and on the spine in Definition of $\mathcal{C}_{p,q}(T)$. A second natural assumption is that the subtrees along the spine of $T \in \mathcal{M}_1(\mathbb{T})$ are iid. It is then possible to completely characterize the corresponding $\mathbb{R}$-tree $\mathcal{T}$, see Proposition 6.4. The construction of $\mathcal{T}$ relates to the quasi-stationary distributions of discrete time subcritical Galton Watson processes, and the link is detailed in Remark 6.6.

Outline of the paper

The paper is organized as follows: Section 2 focuses on the space $\mathfrak{T}_1^\ell$ of probability measured real trees. In this context, a preliminary version of Theorem 1.8, Proposition 2.1, is stated and proved. Section 3 introduces the infinite trees that are needed to treat the case of $\mathfrak{T}^\ell$ and finally prove Theorem 1.8. The proof of our main results, Theorems 1.12 and Theorem 1.17, then follow in Section 4. Motivated by the correspondence between self-similar discrete trees and self-similar measured $\mathbb{R}$-trees established in Theorem 1.17, we provide examples of self-similar measured $\mathbb{R}$-trees in Section 5. Further properties of these trees may be obtained under additional assumptions on the associated discrete trees, and Section 6 discusses two of them: translation invariance, and iid subtrees. An appendix recalls the necessary notions on $\mathbb{R}$-trees.

Acknowledgments

We are grateful to Itai Benjamini for asking us a question which motivated this study. We are also deeply indebted to Ohad Feldheim; example 1 in Section 5 was found following discussions with him before the general statement of Theorem 1.17 was clear to us. Finally, we thank Grégory Miermont and Ron Peled for useful discussions.

2 Trees, partially exchangeable arrays and exchangeable partial orders

In this section, we will prove the following result, which is a weaker version of Theorem 1.8:

Proposition 2.1. *The topology $\mathcal{O}$ coincides with the GHP topology on $\mathfrak{T}_1^\ell$. Furthermore, the space $\mathfrak{T}_1^\ell$ is compact.*

For this we rely on two different representations of random rooted measured trees, one in terms of its distance matrix distribution defined in Section A, the other in terms of an exchangeable partial order on $\mathbb{N}$. We recall that given a tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$, its *distance matrix distribution* $\eta(\mathcal{T})$ is defined as the push-forward of the probability measure $\delta_\rho \otimes \mu^{\otimes \mathbb{N}^*}$ by the map

$$(x_i)_{0 \leqslant i \leqslant n} \mapsto d(x_i, x_j)_{0 \leqslant i,j \leqslant n}.$$

Hence, $\eta(\mathcal{T})$ is a probability measure on $[0,1]^{\mathbb{N}\times\mathbb{N}}$.

Denote by $\mathcal{P}$ the space of exchangeable random partial orders on $\mathbb{N}$, which is a closed subspace of the space of probability measures on $\{0,1\}^{\mathbb{N}\times\mathbb{N}}$, and hence compact. Every tree $\mathcal{T} \in \mathfrak{T}_1^\ell$ then defines an element $\varphi(\mathcal{T}) = \mathcal{D}(\mathcal{T}, \infty)$ in $\mathcal{P}$ as follows: Let $(X_1, S_1), (X_2, S_2), \dots$ be an iid sequence of random variables in $\mathcal{V} \times \{0,1\}$ with law $\ell_\mathcal{T} \otimes \delta_1 + (\mu - \ell_\mathcal{T}) \otimes \delta_0$. This means that $X_1, X_2, \dots$ are iid according to μ and $S_i = 1$ if X_i was drawn according to $\ell_\mathcal{T}$ and $S_i = 0$ otherwise. Set $(X_0, S_0) = (\rho, 1)$. We then define the random partial order $\varphi(\mathcal{T}) = \preceq_\mathcal{T}^\mathbb{N}$ on $\mathbb{N}$ by

$$i \preceq_\mathcal{T}^\mathbb{N} j \iff X_i \preceq_\mathcal{T} X_j \text{ and } S_i = 1.$$

We denote by $\mathcal{P}_\mathfrak{T} = \overline{\varphi(\mathfrak{T}_1^\ell)} \subset \mathcal{P}$ the closure of the range of the map φ.

Conversely, to every exchangeable partial order $\preceq$ from $\mathcal{P}$, we can associate a random distance matrix as follows: In order to recover the n first coordinates of the distance matrix, consider a regular conditional distribution of the partial order restricted to $\{1,\dots,n+1\}$, conditioned on the partial order restricted to $\{1,\dots,n\}$. Denote this distribution by $\mathbf{P}_n$. Then we set

$$d_n(i,j) = \mathbf{P}_n(n+1 \preceq i \text{ and } n+1 \npreceq j) + \mathbf{P}_n(n+1 \preceq j \text{ and } n+1 \npreceq i). \tag{2.1}$$

By exchangeability, $d_n(i,j) \equiv d_m(i,j)$ for every $n,m \geqslant \max(i,j)$, whence there exists $d_\preceq(i,j)$, such that $d_n(i,j) = d_\preceq(i,j)$ for $n \geqslant \max(i,j)$. We denote the map associating this random distance matrix to the partial order by $\psi^* : \mathcal{P} \to \mathcal{M}_1([0,1]^{\mathbb{N}\times\mathbb{N}})$. Clearly, for every $\mathcal{T} \in \mathfrak{T}_1^\ell$, $\psi^*(\varphi(\mathcal{T})) = \eta(\mathcal{T})$. We now have the following result:

Proposition 2.2. *Let $\preceq$ in $\mathcal{P}_\mathfrak{T}$, such that $\psi^*(\preceq) = \eta(\mathcal{T})$ for some tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$ whose measure μ has full support. Then $\mathcal{T} \in \mathfrak{T}_1^\ell$ and $\varphi(\mathcal{T})$ equals $\preceq$.*

Proof. Let $X_1, X_2, \dots$ be iid vertices in the tree chosen according to μ and set $X_0 = \rho$. Now couple the partial order $\preceq$, the distance array $d_\preceq$ and the sequence $X_1, X_2, \dots$ such that (2.1) holds and $d_\preceq(i,j) = d(X_i, X_j)$. We now claim that almost surely,

$$\forall i \neq j \in \mathbb{N} : i \preceq j \iff X_i \preceq_\mathcal{T} X_j \text{ and } i \in \mathrm{Sk} = \{i \in \mathbb{N} : \nexists j \neq i : i \preceq j\} \tag{2.2}$$

Let us show how we can prove the proposition from (2.2). For notational convenience, change the index set to $\mathbb{N} \cup \{\star\}$ for some symbol $\star \notin \mathbb{N}$, by mapping 1 to $\star$ and $n+1$ to n for every $n \geqslant 1$ (by exchangeability, we do not change the law of $\preceq$ and of the X_i). Denote by $\mathbf{P}_\mathbb{N}$ a regular conditional distribution of the partial order $\preceq$ conditioned to its restriction to $\mathbb{N}$. We then have by (2.1) and (2.2), for every $i,j \in \mathbb{N}$,

$$\begin{aligned} d(i,j) = d_\preceq(i,j) &= \mathbf{P}_\mathbb{N}(\star \preceq i \text{ and } \star \npreceq j) + \mathbf{P}_\mathbb{N}(\star \preceq j \text{ and } \star \npreceq i) \\ &= \mathbf{P}_\mathbb{N}(X_\star \in [\![X_i, X_j]\!] \text{ and } \star \in \mathrm{Sk}). \end{aligned} \tag{2.3}$$

Now, since the measure μ of the tree $\mathcal{T}$ has full support, the X_i form a dense subset of the vertices. It follows from (2.3) that for every Borel set $A \subset \mathcal{V}$,

$$\mathbf{P}_\mathbb{N}(X_\star \in A \text{ and } \star \in \mathrm{Sk}) = \ell_\mathcal{T}(A).$$

This shows that $\mu \geqslant \ell_\mathcal{T}$ and therefore $\mathcal{T} \in \mathfrak{T}_1^\ell$. By exchangeability, we now have for every $n \geqslant 0$,

$$\mathbf{P}_n(X_{n+1} \in A \text{ and } \star \in \mathrm{Sk}) = \ell_\mathcal{T}(A) \text{ and } \mathbf{P}_n(X_{n+1} \in A \text{ and } \star \notin \mathrm{Sk}) = \mu(A) - \ell_\mathcal{T}(A).$$

Together with (2.2), this shows that the partial order $\preceq$ is exactly $\varphi(\mathcal{T})$ (with $V_0 = \{X_i : i \in \mathrm{Sk}\}$ and $V_1 = \{X_i : i \notin \mathrm{Sk}\}$).

It now remains to prove (2.2). We first note that

$$\begin{aligned} X_i \preceq_\mathcal{T} X_j &\iff d(\rho, X_i) + d(X_i, X_j) = d(\rho, X_j) \\ &\iff \mathbf{P}_\mathbb{N}(\star \preceq i) + \mathbf{P}_\mathbb{N}(\star \preceq j, \star \npreceq i) + \mathbf{P}_\mathbb{N}(\star \preceq i, \star \npreceq j) = \mathbf{P}_\mathbb{N}(\star \preceq j) \\ &\iff \mathbf{P}_\mathbb{N}(\star \preceq j) + 2\mathbf{P}_\mathbb{N}(\star \preceq i, \star \npreceq j) = \mathbf{P}_\mathbb{N}(\star \preceq j) \\ &\iff \mathbf{P}_\mathbb{N}(\star \preceq i, \star \npreceq j) = 0. \end{aligned}$$

This yields the direct implication "$\Rightarrow$" in (2.2), since $i \preceq j$ trivially implies $\mathbf{P}_\mathbb{N}(\star \preceq i, \star \npreceq j) = 0$ and $i \in \mathrm{Sk}$. For the reverse implication, it suffices to prove the statement with $j = 1$ and $i = 2$. Define the event

$$B_i = \{i \in \mathrm{Sk},\, \mathbf{P}_\mathbb{N}(\star \preceq i, \star \npreceq 1) = 0,\, i \npreceq 1\}.$$

We want to show that $\mathbf{P}(B_2) = 0$, which will prove the reverse implication in (2.2). First note that on $B_i \cap B_k$, $i \neq k$, we have $i \npreceq k$ almost surely, since by exchangeability, we have (exchanging i with $\star$),

$$\mathbf{P}(i \preceq k,\, B_i \cap B_k) \leqslant \mathbf{P}_\mathbb{N}(\star \preceq k,\, \star \npreceq 1,\, \mathbf{P}_\mathbb{N}(\star \preceq k, \star \npreceq 1) = 0) = 0.$$

By symmetry, it follows that on the event $B_i \cap B_k$, we have $i \not\preceq k$ and $k \not\preceq i$ almost surely. We now use the fact that the partial order $\preceq$ is in $\mathcal{P}_\mathfrak{T}$, i.e. it is a limit of ancestral relations of real trees. In particular, it satisfies almost surely,

$$i \preceq \star \text{ and } k \preceq \star \implies i \preceq k \text{ or } k \preceq i.$$

Together with the previous reasoning this yields almost surely,

$$B_i \cap B_k \Rightarrow \mathbf{P}_\mathbb{N}(i \preceq \star,\, k \preceq \star) = 0. \tag{2.4}$$

Now set $Y_i = \mathbf{P}_\mathbb{N}(i \preceq \star)\mathbb{1}_{B_i}$. By de Finetti's theorem, the following limit exists almost surely:

$$Y = \lim_{n\to\infty} \frac{1}{n} \sum_{i=2}^{n} Y_i,$$

Furthermore, by (2.4), we have for every n, almost surely,

$$\sum_{i=2}^{n} Y_i \leqslant \mathbf{P}_\mathbb{N}\left(\bigcup_{i=2}^{n} \{i \preceq \star\}\right) \leqslant 1.$$

Hence, $Y = 0$ almost surely, which implies $\mathbf{P}_\mathbb{N}(2 \preceq \star)\mathbb{1}_{B_2} = 0$ almost surely. But on the event B_2, we have $2 \in \mathrm{Sk}$, hence $\mathbf{P}_\mathbb{N}(2 \preceq \star) > 0$ almost surely by de Finetti's theorem. This implies that $\mathbf{P}(B_2) = 0$, which was to be proven. This finishes the proof of (2.2) and therefore of the proposition. □

Lemma 2.3. *The map $\psi^* : \mathcal{P} \to \mathcal{M}_1([0,1]^{\mathbb{N}\times\mathbb{N}})$ is continuous.*

Proof. By definition, it is enough to show that the law of the distance matrices $(d_n(i,j))_{0\leqslant i,j\leqslant n}$ associated to a partial order $\preceq$ are continuous in $\preceq$. But this follows from the fact that the restriction of the partial order $\preceq$ to $\{1,\dots,n+1\}$ is a random element in a finite state space, and so the conditional distribution in (2.1) is continuous in $\preceq$. □

Proof of Proposition 2.1. In order to prove the Proposition, it suffices to show that the map φ is a homeomorphism between $\mathfrak{T}_1^\ell$ and $\mathcal{P}_\mathfrak{T}$. Since $\mathcal{P}_\mathfrak{T}$ is compact and $\mathfrak{T}_1^\ell$ is Hausdorff, it is enough to show that φ is bijective and its inverse φ^{-1} is continuous [Dug66, Theorem XI.2.1].

Let $\preceq$ in $\mathcal{P}_\mathfrak{T}$. By definition, there exists a sequence of trees $\mathcal{T}_n \in \mathfrak{T}_1^\ell$, such that $\varphi(\mathcal{T}_n)$ converges to $\preceq$. By Lemma A.5, the space $\mathfrak{T}_1^\ell$ is precompact in the GHP topology and the measure μ^* of every subsequential limit $\mathcal{T}_*$ has full support. By Lemma 2.3, we have $\psi^*(\preceq) = \eta(\mathcal{T}_*)$, whence the tree $\mathcal{T}_*$ is uniquely defined by Fact A.7 and the sequence $\mathcal{T}_n$ therefore converges to $\mathcal{T}_*$ in the GHP topology. Furthermore, by Proposition 2.2, we have $\mathcal{T}^* \in \mathfrak{T}_1^\ell$. We can therefore define a map $\psi : \mathcal{P}_\mathfrak{T} \to \mathfrak{T}_1^\ell$ by $\psi(\preceq) = \mathcal{T}_*$. Note that this map satisfies $\eta(\psi(\preceq)) = \psi^*(\preceq)$. It is therefore continuous by Lemma 2.3 and Corollary A.9.

Now, by Proposition 2.2, the maps φ and ψ are inverses. We have therefore proven the statement about φ stated in the beginning of the proof and finished the proof of the theorem. □

3 Infinite trees

We now extend the statement of Proposition 2.1 to the space $\mathfrak{T}^\ell$ in order to prove Theorem 1.8.

Proof of Theorem 1.8. We first note that the GHP and the $\mathcal{O}$ topology coincide on the space $\mathfrak{T}_1^\ell$ by Proposition 2.1. The extension to $\mathfrak{T}_f^\ell$ is trivial: It suffices to rescale distances and the measure μ by its total mass, which yields a tree in $\mathfrak{T}_1^\ell$. By the second part of Proposition 2.1, this moreover gives that every limit point of $\mathfrak{T}_f^\ell$ in $\mathfrak{T}_f$ is also contained in $\mathfrak{T}_f^\ell$.

We now treat the whole space $\mathfrak{T}^\ell$. Let $(\mathcal{T}_n = (\mathcal{V}_n, d_n, \rho_n, \mu_n))_{n\geqslant 0}$ be a sequence in $\mathfrak{T}^\ell$. First assume that $\mathcal{T}_n$ converges as $n \to \infty$ to a limit $\mathcal{T}$ in the GHP topology. By Lemma A.10, this

implies that for every $R \geqslant 0$, there exists a sequence $R_n \to R$ as $n \to \infty$, such that the restricted trees $\mathcal{T}_n^{\leq R_n} \to \mathcal{T}^{\leq R}$ as $n \to \infty$. For every $m \geqslant 0$, one can now choose $R = R(m)$, such that the tree $\mathcal{D}(\mathcal{T}_n)^{\leqslant m}$ is a subtree of $\mathcal{D}(\mathcal{T}_n^{\leq R_n})$ with high probability. The convergence of $\mathcal{D}(\mathcal{T}_n)^{\leqslant m}$ now follows from the fact that the GHP convergence of $\mathcal{T}_n^{\leq R_n}$ implies convergence in the $\mathcal{O}$ topology. This shows that the GHP topology is finer than the $\mathcal{O}$ topology (i.e. the map $\mathcal{D}$ is continuous).

For the converse, assume that $\mathcal{D}(\mathcal{T}_n)$ converges to $\mathcal{D}(\mathcal{T})$ as $n \to \infty$. It is enough to show that this implies that the sequence $(\mathcal{T}_n)_{n\geqslant 0}$ is precompact in the GHP topology, since every limit point $\mathcal{T}^*$ satisfies $\mathcal{D}(\mathcal{T}^*) = \mathcal{D}(\mathcal{T})$ by the continuity established above. But since the topology $\mathcal{O}$ has already been shown to separate points in $\mathfrak{T}_f^\ell$, this implies in particular that $\mathcal{D}$ is injective, whence $\mathcal{T}^* = \mathcal{T}$.

In order to show precompactness, let N_n^R be the number of vertices in $\mathcal{D}(\mathcal{T}_n)^{\leqslant 1}$, conditioned on the fact that the set V_0 from the definition of $\mathcal{D}$ does not contain any vertex on the spine at distance $\leqslant R$ of the root. By the convergence of $\mathcal{D}(\mathcal{T}_n)$ to $\mathcal{D}(\mathcal{T})$, the probability of this event is uniformly bounded from below in n, and the sequence of variables $(N_n^R)_{n\geqslant 1}$ is then tight for every R. This implies that the number of vertices in the tree $\mathcal{D}(\mathcal{T}_n^{\leq R})$ is a tight sequence of random variables, which in turn implies that the sequence $\mu(\mathcal{T}_n^{\leq R})$ is uniformly bounded in n. This shows the first assumption for precompactness in Lemma A.11. The second assumption is easily verified by the convergence of $\mathcal{D}(\mathcal{T}_n)$ to $\mathcal{D}(\mathcal{T}) \in \mathfrak{T}$ and the fact that the discretization operation preserves infinite rays. This proves precompactness of the sequence $(\mathcal{T}_n)_{n\geqslant 1}$ in the GHP topology and finishes the proof of the theorem. □

Proof of Lemma 1.9. The space $\mathfrak{T}_1^\ell$ is compact according to Proposition 2.1. In particular, it is closed in $\mathfrak{T}_1$. As already mentioned, the fact that $\mathfrak{T}_f^\ell$ is closed in $\mathfrak{T}_f$ then follows by rescaling distances and the measure μ by its total mass: for $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_f^\ell$, the rescaled tree $(\mathcal{C}_{\mu(\mathcal{V})}(\mathcal{T})) \in \mathfrak{T}_1^\ell$. The case of $\mathfrak{T}^\ell$ is dealt similarly. If $\mathcal{T}_n \in \mathfrak{T}^\ell$ converge to $\mathcal{T} \in \mathfrak{T}$, then, for each $R > 0$ there exists a sequence R_n converging to R such that $\mathcal{T}_n^{\leqslant R_n} \to \mathcal{T}^{\leqslant R}$, with $\mathcal{T}_n^{R_n} \in \mathfrak{T}_f^\ell$ and $\mathcal{T}^R \in \mathfrak{T}_f$. Since $\mathfrak{T}_f^\ell$ is closed in $\mathfrak{T}_f$, we have that $\mathcal{T}_n^{\leqslant R} \in \mathfrak{T}_f^\ell$, and, $(\mathcal{T}^{\leqslant R})_{R\geqslant 0}$ then builds a compatible sequence that defines a tree in $\mathfrak{T}^\ell$. □

4 Rescaling of trees

In this section, we will prove a coupling lemma (Lemma 4.1 below), which, together with Theorem 1.8, will yield Theorems 1.12 and 1.17.

The space of discrete trees $\mathbb{T}$ is naturally embedded into the space $\mathfrak{T}^\ell$ via the following embedding ι: given $T \in \mathbb{T}$, we define $\iota(T) = (\mathcal{V}, d, \rho, \mu)$ as follows:

- The set of vertices is given by $\mathcal{V} = \{(i, x), i \in T \setminus \{\rho\}, x \in [0, 1)\} \cup \{\rho\}$.
- The distance is defined by $d((i, x), (j, y)) = d_T(i, j) - x - y$ if ($i \notin \mathrm{Anc}(j)$ and $j \notin \mathrm{Anc}(i)$) and $d((i, x), (j, y)) = d_T(i, j) - x + y$ if $j \in \mathrm{Anc}(i)$.
- μ is the Lebesgue measure $\ell_{\mathcal{T}}$.

Informally, the tree $\iota(T)$ is defined from T by adding segments of length 1 between the vertices of T, and the element (i, x) of $\mathcal{V}$ is at distance x of $(i, 0)$ on the path $[\![\rho, i]\!]$ to the root.

Lemma 4.1. *Let $(p_n)_{n\geqslant 0}$ and $(q_n)_{n\geqslant 0}$ by sequences of positive numbers such that $p_n \to 0$ and $q_n \to 0$ as $n \to \infty$. Let $T_n \in \mathcal{M}_1(\mathbb{T})$ be a sequence of random rooted trees. Then the following statements are equivalent:*

1. *The sequence $(\mathcal{C}_{p_n,q_n}(T_n))_{n\geqslant 0}$ is tight.*
2. *The sequence $(\mathcal{D}(\mathcal{S}_{p_n,q_n}(\iota(T_n))))_{n\geqslant 0}$ is tight.*
3. *The sequence $(\mathcal{S}_{p_n,q_n}(\iota(T_n)))_{n\geqslant 0}$ is tight.*

In this case, for every $m \geqslant 1$ and $\varepsilon > 0$, there exists an integer n_0 and for every $n \geqslant n_0$ a coupling between $\mathcal{C}_{p_n,q_n}(T_n)$ and $\mathcal{D}(\mathcal{S}_{p_n,q_n}(\iota(T_n)))$, such that

$$\mathbf{P}\left([\mathcal{C}_{p_n,q_n}(T_n)]^{\leqslant m} = [\mathcal{D}(\mathcal{S}_{p_n,q_n}(\iota(T_n)))]^{\leqslant m}\right) > 1 - \varepsilon. \tag{4.1}$$

Proof. Equivalence of the second and the third point follows from Theorem 1.8. We will now show equivalence between the first and second point, and, in passing, prove (4.1). Recall that by definition of the topology of local convergence on $\mathbb{T}$, a sequence of random discrete trees $\widetilde{T}_n$ in $\mathcal{M}_1(\mathbb{T})$ is tight if and only if for every $m \geqslant 1$, the sequence $(\#V(\widetilde{T}_n^{\leqslant m}))_{n\geqslant 1}$ is tight. In order to apply this, we first need to define a suitable coupling between the operations $\mathcal{C}_{p,q}$ and $\mathcal{S}_{p,q}$ and suitable one-dimensional stochastic processes.

Let $T \in \mathbb{T}$. Recall that in the construction of $\mathcal{C}_{p,q}(T)$, the vertices are colored with two colors, say black and white; the black vertices (and the root) are retained, whereas the white vertices are discarded. Given a tree $T \in \mathbb{T}$ we can couple this operation with two infinite sequences $(B_i)_{i\geqslant 1}$ and $(B'_i)_{i\geqslant 1}$ of iid random variables distributed according to the Bernoulli distribution with parameter p and q, respectively: we perform a breadth-first search of the vertices of the tree starting from the root and color the vertex visited at step i according to B_i (1 = black and 0 = white) if it is an off-spine vertex, and according to B'_i if it is a spine vertex. Given an integer $m \geqslant 1$, we add an additional rule: When a black vertex is visited which has exactly $m-1$ black ancestors, then its subtree is subsequently ignored by the algorithm. This algorithm terminates almost surely since $\mathcal{C}_{p,q}(T)$ is locally finite by definition. The restriction of the tree T to the black vertices then exactly gives the tree $\mathcal{C}_{p,q}(T)^{\leqslant m}$.

In order to construct the tree $\mathcal{D}(\mathcal{S}_{p,q}(\iota(T)))^{\leqslant m}$ one can proceed in a similar manner, but using now two Poisson processes $(P_t)_{t\geqslant 0}$ and $(P'_t)_{t\geqslant 0}$ with parameters p and q, respectively. We leave the details to the reader.

For every $M \geqslant 0$, we now couple the sequences $(B_i)_{i\geqslant 1}$ and $(B'_i)_{i\geqslant 1}$ with the Poisson processes $(P_t)_{t\geqslant 0}$ and $(P'_t)_{t\geqslant 0}$ in such a way that with probability $1-O(M(p\vee q))$, for every $i \leqslant M/p$, $B_i = 1$ if and only if $P_i - P_{i-1} = 1$ and for every $i \leqslant M/q$, $B'_i = 1$ if and only if $P'_i - P'_{i-1} = 1$. It is now easy to check that for large M and small p and q, on the event that (either) $\#V(\mathcal{C}_{p,q}(T)^{\leqslant m}) \leqslant M/2$ or $\#V(\mathcal{D}(\mathcal{S}_{p,q}(\iota(T)))^{\leqslant m}) \leqslant M/2$, both trees agree with high probability. This implies that if the first or second statement of the lemma holds, then (4.1) is true, which in turn implies equivalence of the two statements. This finishes the proof of the lemma. □

Proof of Theorem 1.12. We start with the easy direction: Let $\mathcal{T} \in \mathcal{M}_1(\mathfrak{T}_1^\ell)$. Then the family of random rooted trees $T_n = \mathcal{D}(\mathcal{T}, n)$ is compatible since, for $n \geqslant m \geqslant 1$:

$$\mathcal{C}(T_n, m) = \mathcal{C}(\mathcal{D}(\mathcal{T}, n), m) = \mathcal{D}(\mathcal{T}, m) = T_m.$$

Conversely, starting from a compatible family $(T_n)_{n\in\mathbb{N}} \in \mathcal{M}_1(\mathbb{T}_f)$, we have:

$$\mathcal{C}(T_n, \mathrm{Bin}(n, 1/n)) = T_{\mathrm{Bin}(n,1/n)},$$

where $\mathrm{Bin}(n, 1/n)$ an independent Binomial random variable with parameter n and $1/n$. The right-hand side converges towards $T_{\mathrm{Po}(1)}$ as $n \to \infty$. On the left hand side, we first notice that $\mathcal{C}(T_n, \mathrm{Bin}(n, 1/n)) = \mathcal{C}_{1/n+1}(T_n)$, and then observe that Lemma 4.1 ensures the existence of a coupling such that

$$\mathcal{C}_{1/n+1}(T_n) = \mathcal{D}(\mathcal{S}_{1/(n+1)}(\iota(T_n)))$$

holds with high probability as $n \to \infty$. The space $\mathfrak{T}_1^\ell$ being compact by Proposition 2.1, the sequence $(\mathcal{S}_{1/(n+1)}(\iota(T_n)))_{n\geqslant 1} \in \mathcal{M}_1(\mathfrak{T}_1^\ell)$ is then tight. Since the map $\mathcal{D}$ is continuous for the topology $\mathcal{O}$ by Theorem 1.8, every subsequential limit $\mathcal{T}$ then satisfies $T_{\mathrm{Po}(1)} = \mathcal{D}(\mathcal{T})$, and by conditioning, $T_m = \mathcal{D}(\mathcal{T}, m)$. Last, the topology $\mathcal{O}$ separates points by the same Theorem 1.8, and the tree $\mathcal{T}$ is therefore unique. □

Proof of Theorem 1.17. One direction is obvious: Let $\mathcal{T} \in \mathcal{M}_1(\mathfrak{T}^\ell)$ be a (p,q)-self-similar random real tree. Lemma 1.16 then yields

$$\mathcal{C}_{p,q}(\mathcal{D}(\mathcal{T})) = \mathcal{D}(\mathcal{S}_{p,q}(\mathcal{T})) = \mathcal{D}(\mathcal{T}),$$

whence the discrete tree $\mathcal{D}(\mathcal{T})$ is (p,q)-self-similar as well.

For the other direction, let $T \in \mathcal{M}_1(\mathbb{T})$ be a (p,q)-self-similar random rooted tree, i.e. $\mathcal{C}_{p,q}(T) = T$. Iterating this equality yields for each integer $n \geqslant 1$,

$$\mathcal{C}_{p^n,q^n}(T) = \mathcal{C}^n_{p,q}(T) = T. \tag{4.2}$$

In particular, the sequence $(\mathcal{C}_{p^n,q^n}(T))_{n\geqslant 1}$ is tight. Lemma 4.1 now yields that the sequence $(\mathcal{S}_{p^n,q^n}(\iota(T)))_{n\geqslant 1}$ is tight and that every subsequential limit $\mathcal{T}$ satisfies $\mathcal{D}(\mathcal{T}) = T$. Since the topology $\mathcal{O}$ separates points by Theorem 1.8, the tree $\mathcal{T}$ is unique. Furthermore, by the continuity of the operator $\mathcal{S}_{p,q}$, we have

$$\mathcal{S}_{p,q}(\mathcal{T}) = \lim_{n\to\infty} \mathcal{S}_{p^{n+1},q^{n+1}}(\iota(T)) = \lim_{n\to\infty} \mathcal{S}_{p^n,q^n}(\iota(T)) = \mathcal{T},$$

whence the tree $\mathcal{T}$ is (p,q)-self-similar. This proves the theorem. □

5 Examples

In this section, we construct some examples of (p,q)-self-similar $\mathbb{R}$-trees. We do not believe that it is possible to completely characterize this family, similarly to the situation for self-similar real-valued processes.

1. Subordination of a real-valued self-similar process. If $(\mathcal{T}, d, \rho, \mu)$ is a (p,q)-self-similar $\mathbb{R}$-tree, then let $\mathcal{T}_t$ be the subset consisting of the vertices whose most recent ancestor on the spine is at distance at most t from the root. Setting $X(t) = \mu(\mathcal{T}_t) - t$ defines a (semi)-self-similar real-valued non-decreasing process with Hurst exponent $H = \log p/\log q$, i.e.

$$(q^{-H}X(qt), t \geqslant 0) \text{ is equal in law to } (X(t), t \geqslant 0). \tag{5.1}$$

   On the other hand, if we are given such a process $X(t)$, we can construct from it a (p,q)-self-similar $\mathbb{R}$-tree $(\mathcal{T}, d, \rho, \mu)$. write $X_c(t)$ for its continuous part and $X_j(t)$ for its jump part. Let furthermore $(\mathcal{T}', d', \rho', \mu')$ be a rooted measured $\mathbb{R}$-tree, such that μ' is a probability measure. The tree $(\mathcal{T}, d, \rho, \mu)$ is then constructed as follows:

   - $\mu(\mathrm{Spine}(\mathcal{T}) \cap [0,t]) = X_c(t) + t$,
   - For every jump time t of X_j, we attach an independent copy of $\mathcal{T}'$ to the spine at distance t of the root, rescaled by the size of the jump $X_j(t) - X_j(t-)$.

   It is easy to show that the resulting tree is indeed (p,q)-self-similar.

2. One can easily generalize the above construction. For example, instead of attaching independent rescaled copies of the same tree $\mathcal{T}'$ to the spine, one can take a $|\log q|$-stationary process $(\mathcal{T}'(s), s \in \mathbb{R})$ of rooted, probability-measured $\mathbb{R}$-trees ($|\log q|$-stationary means that $(\mathcal{T}(s), s \in \mathbb{R})$ is equal in law to $(\mathcal{T}(s + |\log q|), s \in \mathbb{R})$), and attach a rescaled copy of $\mathcal{T}'(\log t)$ at the point t on the spine. One can also introduce a stronger dependency between the process $X(t)$ and the trees. For example, let $X(t)$ be as above and suppose for simplicity that it is a pure-jump process. Let $\mathcal{R}$ be the set of its *record jumps*, i.e. $r \in \mathcal{R}$ if and only if $X(r) - X(r-) > X(s) - X(s-)$ for all $s < r$. Write $\mathbb{R} = \{\ldots < r_{-1} < r_0 < r_1 < \ldots\}$, with $r_0 \leqslant 1 < r_1$ and let $\mathcal{T}'_n$ be a sequence of iid copies of a probability-measured rooted $\mathbb{R}$-tree. We then construct a (p,q)-self-similar tree as follows: For every t, let $N(t)$ be such that $r_{N(t)} \leqslant t < r_{N(t)+1}$. Then for each jump time t of $X(t)$, add the tree $\mathcal{T}'_{N(t)}$ to the spine, rescaled by the size $X(t) - X(t-)$ of the jump. One readily checks that the resulting process is (p,q)-self-similar.

3. Given a (p,q)-self-similar $\mathbb{R}$-tree $\mathcal{T}$, one can easily construct a whole family of self-similar $\mathbb{R}$-trees: Let $\beta > 0$. First, one can rescale the tree along the spine: define a new $\mathbb{R}$-tree $\mathcal{T}'$ obtained from $\mathcal{T}$ by mapping a point t on the spine to t^β. The mass process of this new tree is $X^\beta(t) = X(t^{1/\beta})$, and therefore $(p^{-1}X^\beta(q^\beta t); t \geqslant 0) \overset{\text{law}}{=} (X^\beta(t); t \geqslant 0)$. Since we have not changed the structure of the subtrees, it follows that the resulting tree is (p, q^β)-self-similar.

   For $\gamma \in \mathbb{R}$, one can also define a new mass process by setting $X^\gamma(t) = \int_0^t s^\gamma dX(s)$, as long as this quantity is finite for some (hence, any) $t > 0$. The tree defined by this process in the canonical way (i.e. by rescaling the subtrees of the spine and the measure μ on the spine) is then (pq^γ, q)-self-similar.

   Finally, one can apply the previous scaling to the continuous part $X_c(t)$ of the mass process only, and scale the jump process $X_j(t)$ instead by setting $X_j^\delta(t) = \sum_{s \leqslant t}(X_j(s) - X_j(s-))^\delta$ for some $\delta > 0$. If δ and γ are such that $pq^\gamma = p^\delta$, then the mass process $X^{\gamma,\delta}(t) = X_c^\gamma(t) + X_j^\delta(t)$, if it exists, defines a (p^δ, q)-self-similar tree in the canonical way.

6 Translation invariant self-similar trees

In this section, we study self-similar trees which are invariant under translation along the spine. For simplicity, we restrict ourselves to one-ended self-similar trees. We denote the corresponding subspaces of $\mathbb{T}$ and $\mathfrak{T}^\ell$ by $\overline{\mathbb{T}}$ and $\overline{\mathfrak{T}}^\ell$, respectively. Throughout the section, $T \in \mathcal{M}_1(\overline{\mathbb{T}})$ will denote a one-ended (p,q)-self-similar tree and $\mathcal{T} = (\mathcal{T}, d, \rho, \mu) \in \mathcal{M}_1(\overline{\mathfrak{T}}^\ell)$ its limiting $\mathbb{R}$-tree from Theorem 1.17. We define a shift operator $S : \mathcal{M}_1(\overline{\mathbb{T}}) \to \mathcal{M}_1(\overline{\mathbb{T}})$ which maps T to the subtree rooted at the vertex on the spine at distance 1 from the root. We say T is *translation invariant*, if $ST = T$. Similarly, for $t \geqslant 0$, we define a shift operator $S_t : \mathcal{M}_1(\overline{\mathfrak{T}}^\ell) \to \mathcal{M}_1(\overline{\mathfrak{T}}^\ell)$ mapping $\mathcal{T}$ to its subtree rooted at the vertex on the spine at distance t from the root. Note that $(S_t)_{t \geqslant 0}$ is a semigroup. We then say that $\mathcal{T}$ is translation invariant if $S_t\mathcal{T} = \mathcal{T}$ for every $t \geqslant 0$. The following proposition says that the two notions are equivalent:

Proposition 6.1. *T is translation invariant if and only if $\mathcal{T}$ is translation invariant.*

Lemma 6.2. *The semigroup $(S_t)_{t \geqslant 0}$ is strongly left-continuous, i.e. for every $\mathcal{T} \in \mathcal{M}_1(\overline{\mathfrak{T}}^\ell)$, the function $t \mapsto S_t(\mathcal{T})$ is left-continuous.*

Proof. It is enough to check that that the shift operator is left-continuous when operating on deterministic trees. Let $\mathcal{T} \in \mathfrak{T}^\ell$. But, using definition A.2 of the d_{GHP} distance, this amounts to prove that for each $r > 0$, $d^c_{GHP}(\mathcal{T}_1^{\leqslant r+\varepsilon}, \mathcal{T}_2^{\leqslant r}) \to 0$ as $\varepsilon \to 0$, where $\mathcal{T}_1 = S_{t-\varepsilon}(\mathcal{T})$ and $\mathcal{T}_2 = S_t(\mathcal{T})$. Within the three terms in the definition A.1 of d^c_{GHP}, we need to verify that only two go to 0 according to Lemma A.6, since $\mathcal{T}_1$ and $\mathcal{T}_2$ belong to $\mathfrak{T}^\ell_f$: the distance between the roots, and the Prokhorov distance. We then need an embedding ϕ: since the two trees $\mathcal{T}_1$ and $\mathcal{T}_2$ are already embedded in a common metric space, it is natural to use this embedding. Now, the distance between the roots satisfies $d(\rho_1, \rho_2) = \varepsilon$ by definition, and the Prokhorov distance between measures μ_1 and μ_2 has: $d^c_P(\mu_1, \mu_2) \leqslant \mu(\mathcal{T}^{<t} \backslash \mathcal{T}^{<t-\varepsilon})$, where $(\mathcal{T})^{<r}$ is defined as the restriction of $\mathcal{T}$ to the *open* ball of radius r. Since the difference $\mathcal{T}^{<t} \backslash \mathcal{T}^{<t-\varepsilon}$ decreases to the null set, its measure goes to 0 as $\varepsilon \to 0$ and the lemma is proved. □

Proof of Proposition 6.1. By Theorem 1.17, the tree T is obtained from $\mathcal{T}$ by sampling vertices according to a Poisson process on $\mathcal{T}$ with intensity μ. By definition of the sampling procedure, the vertices on the spine are those which are sampled according to ℓ_{Spine}, where $\ell_{\text{Spine}} = \ell_{\mathcal{T}}|_{\text{Spine}(\mathcal{T})}$ is the restriction of the length measure $\ell_{\mathcal{T}}$ to the spine. It follows that

$$S^n T = \mathcal{D}(S_{X_1 + \ldots + X_n}\mathcal{T}), \tag{6.1}$$

where $X_1, \ldots, X_n$ are independent and exponentially distributed random variables with parameter 1, independent from $\mathcal{T}$. Setting $n = 1$ in (6.1) directly yields the "if" statement of the proposition.

Now suppose that T is translation invariant. Then, for every $n \geqslant 0$ and $k \geqslant 0$, we have

$$\begin{aligned}
\mathcal{D}(\mathcal{T}) &= T && \text{by hypothesis on } T \text{ and } \mathcal{T} \\
&= \mathcal{C}^k_{p,q} T && \text{by the self-similarity of } T \\
&= \mathcal{C}^k_{p,q} S^n T && \text{by translation invariance of } T \\
&= \mathcal{C}^k_{p,q} \mathcal{D}(S_{X_1+\ldots+X_n}\mathcal{T}) && \text{by (6.1)} \\
&= \mathcal{D}(\mathcal{S}^k_{p,q} S_{X_1+\ldots+X_n}\mathcal{T}) && \text{by Lemma 1.16} \\
&= \mathcal{D}(S_{q^k X_1+\ldots+q^k X_n}\mathcal{S}^k_{p,q}\mathcal{T}) && \text{by definition} \\
&= \mathcal{D}(S_{q^k X_1+\ldots+q^k X_n}\mathcal{T}) && \text{by the self-similarity of } \mathcal{T}.
\end{aligned}$$

Since $\mathcal{D}$ is injective by Theorem 1.8, this implies for every $n \geqslant 0$ and $k \geqslant 0$,

$$\mathcal{T} = S_{q^k X_1+\ldots+q^k X_n}\mathcal{T}.$$

By the law of large numbers, for any $t \geqslant 0$, we can now let $k, n \to \infty$ in such a way that $q^k X_1 + \ldots + q^k X_n$ converges from the left to t almost surely. By Lemma 6.2, this yields the statement. □

We now show that for translation invariant trees, the range of possible values for p and q is restricted. For this, recall the definition of the mass process $X(t)$ from Section 5, and note that $X(t)$ is an real-valued, increasing, semi-self-similar process with Hurst exponent $H = \log p/\log q$. Further recall that $X_c(t)$ and $X_j(t)$ denote, respectively, the continuous part and the jump part of $X(t)$, and note that $X_c(t) = \mu(\mathrm{Spine}(\mathcal{T}) \cap [0,t]) - t$.

Proposition 6.3. *Suppose $\mathcal{T} \in \mathcal{M}_1(\overline{\mathfrak{T}}^\ell)$ is (p,q)-self-similar, translation invariant and non-degenerate (i.e. it is not isometric to $\mathbb{R}_+$). Then $q \geqslant p$. Moreover, if $q > p$, then $X_c \equiv 0$ almost surely, and if $q = p$, then $X_j \equiv 0$ almost surely and $X_c(t) = X(1)t$ for every $t \geqslant 0$.*

Proof. Define the mass process $X(t)$ as in Section 5, such that $X(t)$ is an real-valued, increasing, semi-self-similar process with Hurst exponent $H = \log p/\log q$. The proposition then follows from Theorem 3.3 in [Ver85]. That article actually studied processes for which (5.1) holds for *every* $p, q > 0$ with $\log p/\log q = H$, but the proof can be adapted. □

We now want to study the (p,q)-self-similar discrete trees T, which are translation invariant *and* for which the subtrees along the spine are independent (hence, iid).

Proposition 6.4. *The subtrees of T along the spine are iid if and only if $q \geqslant p$ and $\mathcal{T}$ is constructed as follows:*

- *If $q > p$, then the point process with atoms $(t, \mathcal{T}_t)$, where $\mathcal{T}_t$ is the subtree at the spine vertex at distance t from the root, is a Poisson point process with intensity $dt \otimes \nu(d\mathcal{T})$, with the measure ν decomposing as follows: There exists a measurable family of probability measures σ_x on the space of rooted, probability-measures $\mathbb{R}$-tree, such that $\sigma_x = \sigma_{px}$ for every x, and a measure λ on $(0,\infty)$, such that $\lambda([x,\infty)) = x^{-\log q/\log p} f(\log x)$ for some $|\log p|$-periodic function f, such that ν decomposes as the semi-direct product $\nu = \lambda(dx)\sigma_x^x$, where σ_x^x is the image of the measure σ_x under the map $\mathcal{T} \mapsto x\mathcal{T}$.*
- *If $q = p$, then $\mathcal{T}$ is the line $\mathbb{R}_+$ with a (deterministic) multiple of Lebesgue measure.*

Proof. By Proposition 6.3, we either have $q > p$ and the mass process $X(t)$ is a pure-jump process, or $q = p$ and the mass process is continuous with $X(t) = tX(1)$. The statement in the case $q = p$ now follows from the fact that the subtrees of T are independent if and only if $X(1)$ is a deterministic constant.

Now consider the case $q > p$. By Proposition 6.3, the mass process $X(t)$ is then a pure-jump process, i.e. the restriction of μ to the spine equals the length measure. Denote by T^i for every

$i \geqslant 0$ the off-spine subtree of the i-th vertex on the spine of T. Furthermore, for $s \leqslant t$, denote by $\mathcal{T}^{s,t}$ the *concatenation* of the off-spine subtrees of the tree $\mathcal{T}$ rooted at vertices v on the spine with $s \leqslant d(0,v) < t$. Here, the *concatenation* of a collection of rooted trees is defined to be the rooted tree obtained from the disjoint union of the trees by identification of the roots. By the definition of the discretization operation $\mathcal{D}$, we then have

$$(T^0, T^1, T^2, \ldots) \stackrel{\text{law}}{=} (\mathcal{D}(\mathcal{T}^{\xi_0,\xi_1}), \mathcal{D}(\mathcal{T}^{\xi_1,\xi_2}), \ldots), \tag{6.2}$$

where $\xi_n = \sum_{k=0}^{n} X_k$, with $X_1, X_2, \ldots$ a sequence of independent and exponentially distributed random variables with parameter 1, independent from $\mathcal{T}$. Equation (6.2) now shows that if the point process with atoms $(t, \mathcal{T}_t)$ is translation invariant Poisson process, then the trees $(T^i)_{i\geqslant 0}$ are iid.

Now assume that the trees $(T^i)_{i\geqslant 0}$ are iid. Then as in the proof of Proposition 6.1, we have for every $k \in \mathbb{Z}$, by (6.2) and the self-similarity of T and $\mathcal{T}$,

$$(T^0, T^1, T^2, \ldots) \stackrel{\text{law}}{=} (\mathcal{D}(\mathcal{T}^{q^k\xi_0,q^k\xi_1}), \mathcal{D}(\mathcal{T}^{q^k\xi_1,q^k\xi_2}), \ldots).$$

As in the proof of Proposition 6.1, relying on Lemma 6.2, this implies that the collection of trees $(\mathcal{T}^{t_0,t_1}, \mathcal{T}^{t_1,t_2}, \ldots)$ is independent for every $0 = t_0 < t_1 < t_2 < \ldots$. This in turn implies that the point process $(t, \mathcal{T}_t)$ is a Poisson process, and by translation invariance, its intensity measure is of the form $\sigma = dt \otimes \nu$ for some measure ν. By the (p,q)-self-similarity, $\nu = q^{-1}\nu^p$, where ν^p is the image of the measure under the map $\mathcal{T} \mapsto p\mathcal{T}$. Disintegrating the measure ν with respect to the mass of the tree yields the decomposition stated in the theorem. □

The following corollary follows from Proposition 6.4 and standard properties of Poisson processes. Recall that we define the *concatenation* of a collection of rooted trees to be the rooted tree obtained from the disjoint union of the trees by identification of the roots.

Corollary 6.5. *The subtrees of T along the spine are iid if and only if $q \geqslant p$ and the subtrees of T are independent copies of a tree T^0, where T^0 is as follows: There exists a constant $c > 0$, as well as a measurable family of probability measures σ_x and a measure λ as in the statement of Proposition 6.4, such that*

- *T^0 is the concatenation of copies of $\mathrm{Geo}(1/(1+c))$-distributed number of independent copies of a tree T', where*
- *$\mathbf{E}[F(T')] = c^{-1}\int_0^\infty \lambda(dx)\int \sigma_x(d\mathcal{T})\mathbf{E}[F(\mathcal{D}(\mathcal{T}, \mathrm{Poi}(x)))]$ for every bounded measurable function F with $F(\rho) = 0$ (the constant c equals $\int_0^\infty (1-e^{-x})\lambda(dx)$).*

Remark 6.6. The number of edges in the tree T' that appears in the statement of Corollary 6.5 follows the law

$$\mathbf{P}[N(T') = k] = c^{-1}\int_0^\infty \mathbf{P}(\mathrm{Poi}(x) = k)\lambda(dx)\,\forall k \geqslant 1, \quad \text{with } c = \int_0^\infty (1-e^{-x})\lambda(dx).$$

The laws of this form with λ as in the statement of Proposition 6.4 are exactly the quasi-stationary distributions with eigenvalue $\alpha = \log q/\log p$ of the Markov chain $(Z_{\log p^{-1}n}; n = 0, 1, 2, \ldots)$, where $(Z_t; t \geqslant 0)$ is the standard pure death process (i.e., $Z_{t-} \to Z_t - 1$ with rate Z_{t-}) killed at 0 [HM].

A Measured real trees

There are several equivalent definitions of an $\mathbb{R}$-tree. Maybe the most intuitive one is the following taken from [ADH13]:

Definition A.1. An $\mathbb{R}$-tree is a metric space $(\mathcal{V}, d)$ with the following properties:

1. It is *geodesically linear*, i.e. for every $x, y \in \mathcal{V}$, there is a unique isometry $f_{x,y} : [0, d(x,y)] \to \mathcal{V}$ such that $f_{x,y}(0) = x$ and $f_{x,y}(d(x,y)) = y$.

2. It is "without loops", i.e. for every $x, y \in \mathcal{V}$, if r and q are continuous injective maps from $[0,1]$ to $\mathcal{V}$ such that $q(0) = x$ and $q(1) = y$, and $r(0) = x$ and $r(1) = y$, then $q([0,1]) = r([0,1])$.

If $x, y \in \mathcal{V}$, we use the notation $[\![x, y]\!]$ to denote the range of the geodesic $f_{x,y}$.

An equivalent definition (Theorem 3.40 and Lemma 3.12 in [Eva08]) is stated in the following lemma. The inequality appearing there is also called the *four-point condition*:

Fact A.2. *A metric space $(\mathcal{V}, d)$ is an $\mathbb{R}$-tree if and only if is is (topologically) connected and for every $x, y, z, t \in \mathcal{V}$,*

$$d(x,y) + d(z,t) \leqslant \max\{d(x,z) + d(y,t),\, d(y,z) + d(x,t)\}.$$

A rooted $\mathbb{R}$-tree $\mathcal{T} = (\mathcal{V}, d, \rho)$ is an $\mathbb{R}$-tree $(\mathcal{V}, d)$ together with a distinguished vertex $\rho \in \mathcal{V}$ called the *root*.

We now define *measured* $\mathbb{R}$-trees. We will need to consider $\mathbb{R}$-trees endowed with an infinite measure; these have been considered in [ADH13] and we follow their treatment. A rooted measured $\mathbb{R}$-tree is then by definition a quadruple $\mathcal{T} = (\mathcal{V}, d, \rho, \mu)$ where $(\mathcal{V}, d, \rho)$ is a rooted $\mathbb{R}$-tree and μ is a Borel measure on $\mathcal{V}$. Two measured rooted $\mathbb{R}$-trees are said to be *equivalent* if there exists a root- and measure-preserving isometry between them. As usual, we will identify a tree with its equivalence class.

Definition A.3. We define $\widetilde{\mathfrak{T}}$ the space of equivalence classes of rooted measured $\mathbb{R}$-trees $\mathcal{T} = (\mathcal{V}, d, \rho, \mu)$, where

1. the metric space $(\mathcal{V}, d)$ is complete and locally compact,

2. the measure μ is boundedly finite, i.e. $\mu(A) < \infty$ for every bounded Borel set A.

We further denote by $\mathfrak{T}_f$ the subspace of *compact* trees (in particular, the measure μ is then finite), and by $\mathfrak{T}_1$ the subspace of $\mathfrak{T}_f$ where μ is a probability measure.

We now want to define a metric on the space $\widetilde{\mathfrak{T}}$, which will be called the Gromov–Hausdorff–Prokhorov (GHP) metric d_{GHP}. We first recall the definition of the Hausdorff (pseudo-)metric between two subsets of a metric space (Z, d^Z):

$$d_H^Z(A,B) = \inf\{\varepsilon > 0 : A \subset B^\varepsilon \text{ and } B \subset A^\varepsilon\}, \quad A, B \subset Z,$$

where for $A \subset Z$ we define

$$A^\varepsilon = \{x \in Z : d^Z(x, A) < \varepsilon\}, \quad \text{where } d^Z(x,A) = \inf_{y \in A} d^Z(x,y).$$

Furthermore, we recall the definition of the Prokhorov metric on the space $\mathcal{M}_f(Z)$ of finite Borel measures on Z:

$$d_P^Z(\mu,\nu) = \inf\{\varepsilon > 0 : \mu(F) \leqslant \nu(F^\varepsilon) + \varepsilon \text{ and } \nu(F) \leqslant \mu(F^\varepsilon) + \varepsilon \text{ for all closed } F \subset Z\}.$$

We can now define the GHP metric on the space $\mathfrak{T}_f$ of *compact* trees. For $\mathcal{T} = (\mathcal{V}, d, \rho, \mu), \mathcal{T}' = (\mathcal{V}', d', \rho', \mu') \in \mathfrak{T}_f$, set

$$d_{GHP}^c(\mathcal{T}, \mathcal{T}') = \inf_{\varphi, \varphi', Z} \left[d^Z(\varphi(\rho), \varphi'(\rho')) + d_H^Z(\varphi(\mathcal{V}), \varphi'(\mathcal{V}')) + d_P^Z(\varphi_* \mu, \varphi'_* \mu') \right]. \tag{A.1}$$

Here, the infimum is taken over all isometric embeddings $\varphi : \mathcal{V}_1 \hookrightarrow Z$, $\varphi' : \mathcal{V}_2 \hookrightarrow Z$ into some common complete separable metric space (Z, d^Z) and $\varphi_* \mu$ is the push-forward of the measure μ by the map φ.

We now extend this metric to the space $\widetilde{\mathfrak{T}}$. For a tree $\mathcal{T} \in \widetilde{\mathfrak{T}}$ and $r \geqslant 0$, denote by $\mathcal{T}^{\leqslant r}$ its restriction to the closed ball of radius r around the root (i.e. restriction of the underlying metric space as well as of the measure). We then define the GHP metric d_{GHP} on $\widetilde{\mathfrak{T}}$ by

$$d_{GHP}(\mathcal{T}_1, \mathcal{T}_2) = \int_0^\infty e^{-r} \left(1 \wedge d^c_{GHP}\left(\mathcal{T}_1^{\leqslant r}, \mathcal{T}_2^{\leqslant r}\right)\right) dr. \tag{A.2}$$

The following facts are, respectively, Corollary 3.2 and Proposition 2.10 in [ADH13].

Fact A.4. *1. The space $(\widetilde{\mathfrak{T}}, d_{GHP})$ is a complete separable metric space.*

2. The metrics d_{GHP} and d^c_{GHP} induce the same topology on $\mathfrak{T}_f$ and $\mathfrak{T}_1$.

The *length measure* of a rooted real tree $\mathcal{T} = (\mathcal{V}, d, \rho)$ is by definition the unique σ-finite measure $\ell_{\mathcal{T}}$ on $\mathcal{V}$, such that [Eva08, Section 4.3.5]

$$\forall a, b \in \mathcal{V} : \ell_{\mathcal{T}}(\llbracket a, b \rrbracket) = d(a, b) \quad \text{and} \quad \ell_{\mathcal{T}}(\mathcal{V} \backslash \mathcal{V}^\circ) = 0, \text{ where } \mathcal{V}^\circ = \bigcup_{x \in \mathcal{V}} \llbracket \rho, x \llbracket.$$

We then define the spaces $\widetilde{\mathfrak{T}}^\ell$, $\mathfrak{T}_f^\ell$ and $\mathfrak{T}_1^\ell$ as follows:

$$\widetilde{\mathfrak{T}}^\ell = \{\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \widetilde{\mathfrak{T}} : \mu \geqslant \ell_{\mathcal{T}}\}, \quad \mathfrak{T}_f^\ell = \mathfrak{T}_f \cap \widetilde{\mathfrak{T}}^\ell, \quad \mathfrak{T}_1^\ell = \mathfrak{T}_1 \cap \widetilde{\mathfrak{T}}^\ell.$$

Lemma A.5. *1. A family $\mathfrak{S} \subset \widetilde{\mathfrak{T}}^\ell$ is precompact in $\widetilde{\mathfrak{T}}$ if and only if for every $r \geqslant 0$,*

$$\sup_{\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{S}} \mu(\mathcal{T}^{\leqslant r}) < \infty.$$

2. The space $\mathfrak{T}_1^\ell$ is precompact in $\mathfrak{T}_1$. Furthermore, the measure μ of every limit point $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \widetilde{\mathfrak{T}}$ has full support (i.e. $\operatorname{supp} \mu = \mathcal{V}$).

Proof. The first part and the first statement of the second part follow from combining Theorems 2.11 or 2.6 in [ADH13] and Lemma 4.37 in [Eva08]. For the second statement, let $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$ be a limit of a sequence $\mathcal{T}_n = (\mathcal{V}_n, d_n, \rho_n, \mu_n) \in \mathfrak{T}_1$. Combining the proofs of Lemmas 5.8 and A.1 in [GPW09], we can assume the trees $(\mathcal{T}_n)_{n \geqslant 1}$ and $\mathcal{T}$ are embedded in a single metric space Z. For $x \in Z$, denote by $D_\varepsilon(x)$ the closed ball of radius ε around x. Note that for every point $y \in \mathcal{V}_n$, we have $\mu_n(D_\varepsilon(y)) \geqslant \ell_{\mathcal{T}_n}(D_\varepsilon(y)) \geqslant \varepsilon$ as soon as $\varepsilon < \operatorname{diam}(\mathcal{T}_n)/2$, because the point y is then connected to the boundary of the ball by a path of length at least ε. Fix $\varepsilon > 0$, and $x \in \mathcal{V}$. For n large enough, there exists, $y \in \mathcal{V}_n$ such that $d(x, y) \leqslant \varepsilon/2$, and, therefore, $D_{\varepsilon/2}(y) \subseteq D_\varepsilon(x)$. By the Portmanteau theorem, we have

$$\mu(D_\varepsilon(x)) \geqslant \lim_{n \to \infty} \mu_n(D_\varepsilon(x)) \geqslant \lim_{n \to \infty} \mu_n(D_{\varepsilon/2}(y)) \geqslant \frac{1}{2} \limsup_{n \to \infty} (\varepsilon \wedge \operatorname{diam}(\mathcal{T}_n)).$$

We then distinguish according to the value of $\limsup \operatorname{diam}(\mathcal{T}_n)/2$. If it is positive, then the lemma is proved since ε can be chosen arbitrarily small. If it is null, then we have $\mu(D_\varepsilon(x)) = 1$, which also implies that x is in the $\operatorname{Supp}(\mu)$. □

On the previously defined spaces, the Hausdorff distance appearing in the GHP metric is actually unnecessary: Define the Gromov–Prokhorov (GP) metric d^c_{GP} on $\mathfrak{T}_f^\ell$ by

$$d^c_{GP}(\mathcal{T}, \mathcal{T}') = \inf_{\varphi, \varphi', Z} \left[d^Z(\varphi(\rho), \varphi'(\rho')) + d_P^Z(\varphi_* \mu, \varphi'_* \mu')\right].$$

We then have the following lemma:

Lemma A.6. *We have $d^c_{GP} \leqslant d^c_{GHP} \leqslant 3 d^c_{GP}$ on $\mathfrak{T}_f^\ell$.*

Proof. The first inequality is immediate. For the second inequality, let $\mathcal{T} = (\mathcal{V}, d, \rho, \mu), \mathcal{T}' = (\mathcal{V}', d', \rho', \mu') \in \mathfrak{T}_f$ and $\varepsilon > d^c_{GP}(\mathcal{T}, \mathcal{T}')$. Suppose without loss of generality that the trees $\mathcal{T}$ and $\mathcal{T}'$ are subspaces of the same complete separable metric space (Z, d), such that $d(\rho, \rho') + d_P(\mu, \mu') < \varepsilon$, in particular, $d_P(\mu, \mu') < \varepsilon$. We claim that $d_H(\mathcal{V}, \mathcal{V}') \leqslant 2\varepsilon$. To see this, first observe that $\mathcal{V}$ is closed and $\mathcal{V}'^{\varepsilon}$ is open in Z, whence $\mathcal{V} \backslash \mathcal{V}'^{\varepsilon}$ is closed. The definition of d_P then gives

$$\mu(\mathcal{V} \backslash \mathcal{V}'^{\varepsilon}) < \mu'\left((\mathcal{V} \backslash \mathcal{V}'^{\varepsilon})^{\varepsilon}\right) + \varepsilon = \varepsilon,$$

where the last equality follows from the fact that $\mu'\left((\mathcal{V} \backslash \mathcal{V}'^{\varepsilon})^{\varepsilon}\right) \leqslant \mu'(Z \backslash \mathcal{V}') = 0$. Now suppose that there exists $x \in \mathcal{V}$, such that $d(x, \mathcal{V}'^{\varepsilon}) > \varepsilon$. Since $\mathcal{V}$ is connected, there exists a point $y \in \mathcal{V}$, such that $\varepsilon < d(x, y) < d(x, \mathcal{V}'^{\varepsilon})$. We then must have $[\![x, y]\!] \cap \mathcal{V}'^{\varepsilon} = \varnothing$, otherwise there would be a point $z \in \mathcal{V}'^{\varepsilon}$ with $d(x, z) \leqslant d(x, y)$, whence we would have $d(x, \mathcal{V}'^{\varepsilon}) \leqslant d(x, y)$. Hence, the whole segment $[\![x, y]\!]$ lies in $\mathcal{V} \backslash \mathcal{V}'^{\varepsilon}$, and by the definition of the length measure $\ell_{\mathcal{T}}$ and the space $\mathfrak{T}^{\ell}_f$, we then have

$$\mu(\mathcal{V} \backslash \mathcal{V}'^{\varepsilon}) \geqslant \ell_{\mathcal{T}}(\mathcal{V} \backslash \mathcal{V}'^{\varepsilon}) \geqslant \ell_{\mathcal{T}}([\![x, y]\!]) = d(x, y) > \varepsilon.$$

But this is a contradiction to the previous display equation.

To conclude, we have $d(x, \mathcal{V}'^{\varepsilon}) \leqslant \varepsilon$ for every $x \in \mathcal{V}$, whence $\mathcal{V} \subset \mathcal{V}'^{2\varepsilon+\delta}$ for every $\delta > 0$. Similarly, we have $\mathcal{V}' \subset \mathcal{V}^{2\varepsilon+\delta}$ for every $\delta > 0$. This implies that $d_H(\mathcal{V}, \mathcal{V}') \leqslant 2\varepsilon$. By the definition of d^c_{GHP}, this in turn yields $d^c_{GHP}(\mathcal{T}, \mathcal{T}') \leqslant 3\varepsilon$. Since $\varepsilon > d^c_{GP}(\mathcal{T}, \mathcal{T}')$ was chosen arbitrarily, this finishes the proof. □

The metric d^c_{GP} gives rise to a topology on $\mathfrak{T}_f$ called the Gromov–Prokhorov topology. We now give an equivalent definition of this topology on the space $\mathfrak{T}_1$. We follow [GPW09] (which is heavily influenced by Chapter $3\frac{1}{2}$ in [Gro07]). They consider the case of unrooted trees (or metric spaces), but the results can be easily generalized to the rooted case, for example by identifying the rooted, probability measured tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$ with the unrooted, probability measured tree $(\mathcal{V}, d, \frac{1}{3}\mu + \frac{2}{3}\delta_\rho)$. Given a tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$, we define a *distance matrix distribution* $\eta(\mathcal{T})$, i.e. a probability measure on $[0, 1]^{\mathbb{N} \times \mathbb{N}}$, as the push-forward of the probability measure $\delta_\rho \otimes \mu^{\otimes \mathbb{N}^*}$ by the map

$$(x_i)_{0 \leqslant i \leqslant n} \mapsto d(x_i, x_j)_{0 \leqslant i,j \leqslant n}.$$

Proposition 2.6 in [GPW09] now implies the following:

Fact A.7. *A tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}_1$ whose measure μ has full support is uniquely determined by its distance-matrix distribution $\eta(\mathcal{T})$.*

Furthermore, by Theorem 5 in [GPW09], we have:

Fact A.8. *Let $\mathcal{T}, \mathcal{T}_1, \mathcal{T}_2, \ldots \in \mathfrak{T}_1$. Then $d^c_{GP}(\mathcal{T}_n, \mathcal{T}) \to 0$ if and only if $\eta(\mathcal{T}_n) \to \eta(\mathcal{T})$ as $n \to \infty$ (the latter convergence being in the sense of convergence of finite-dimensional distributions).*

Fact A.8 and Lemma A.6 yield the following corollary:

Corollary A.9. *Let $\mathcal{T}, \mathcal{T}_1, \mathcal{T}_2, \ldots \in \mathfrak{T}^{\ell}_1$. Then $d^c_{GHP}(\mathcal{T}_n, \mathcal{T}) \to 0$ if and only if $\eta(\mathcal{T}_n) \to \eta(\mathcal{T})$ as $n \to \infty$.*

We now define the subspaces $\mathfrak{T} \subset \widetilde{\mathfrak{T}}$ and $\mathfrak{T}^{\ell} \subset \widetilde{\mathfrak{T}}^{\ell}$ of trees having only a finite number of infinite ends. We denote the union of these ends the *spine*, denoted by $\mathrm{Spine}(\mathcal{T}) \subset \mathcal{V}$ for a tree $\mathcal{T} = (\mathcal{V}, d, \rho, \mu) \in \mathfrak{T}$. It is possible to construct a complete separable metric on the space $\mathfrak{T}$ which generates the GHP topology: For a tree $\mathcal{T} \in \mathfrak{T}$ and $0 \leqslant s \leqslant r$, denote by $N_{s,r}(\mathcal{T})$ the number of vertices at distance s of the root in $\mathcal{T}$ that have a descendant at distance r of the root. We then define the metric

$$d'_{GHP}(\mathcal{T}, \mathcal{T}') = d_{GHP}(\mathcal{T}, \mathcal{T}') + \int_0^\infty \int_0^\infty e^{-s-r} \, \mathbb{1}_{(N_{s,s+r}(\mathcal{T}) = N_{s,s+r}(\mathcal{T}'))} \, ds \, dr.$$

One can check that this metric generates the GHP topology and that the space $\mathfrak{T}$ is complete under this metric. A similar metric can be defined to show that the space $\mathbb{T}$ is topologically complete.

In the space $\mathfrak{T}$, instead working with the restriction $\mathcal{T}^{\leqslant r}$, it is convenient to work with a different restriction: For a tree $\mathcal{T} \in \mathfrak{T}$, we define the tree $\mathcal{T}^{\preceq r}$ to be the tree restricted to the vertices whose most recent common ancestor on the spine is at distance at most r from the root. We then have the following characterization of convergence in $\mathfrak{T}$ whose proof we omit:

Lemma A.10. *A sequence $\mathcal{T}_1, \mathcal{T}_2, \ldots \in \mathfrak{T}$ converges to $\mathcal{T} \in \mathfrak{T}$ if and only if for every $R \geqslant 0$ there exists a sequence $R_n \to R$, such that*

$$\mathcal{T}_n^{\preceq R_n} \to \mathcal{T}^{\preceq R} \text{ in } \mathfrak{T}_f.$$

The precompactness criterion in the space $\mathfrak{T}$ now reads as follows:

Lemma A.11. *A family $\mathfrak{S} \subset \mathfrak{T}^\ell$ is precompact in $\mathfrak{T}$ if and only if the following two properties hold:*

1. *For every $r \geqslant 0$,*
$$\sup_{\mathcal{T}=(\mathcal{V},d,\rho,\mu)\in\mathfrak{S}} \mu(\mathcal{T}^{\preceq r}) < \infty.$$
2. *There exists a constant $C \geqslant 0$, such that for every $s \geqslant 0$, for all $r \geqslant r_0 = r_0(s)$,*
$$\sup_{\mathcal{T}=(\mathcal{V},d,\rho,\mu)\in\mathfrak{S}} N_{s,r}(\mathcal{T}) < C.$$

References

[ADH13] Romain Abraham, Jean-François Delmas, and Patrick Hoscheit. A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electronic Journal of Probability*, 18(no.14):1–21, January 2013.

[Ald93] David Aldous. The Continuum Random Tree III. *The Annals of Probability*, 21(1):248–289, January 1993.

[AP98] David Aldous and Jim Pitman. Tree-valued markov chains derived from galton-watson processes. In *Annales de l'Institut Henri Poincare (B) Probability and Statistics*, volume 34, pages 637–686. Elsevier, 1998.

[Bil09] Patrick Billingsley. *Convergence of probability measures*, volume 493. John Wiley & Sons, 2009.

[Drm09] Michael Drmota. *Random trees*. SpringerWienNewYork, Vienna, 2009.

[Dug66] James Dugundji. *Topology*. Allyn and Bacon, Inc., Boston, 1966.

[DW07] Thomas Duquesne and Matthias Winkel. Growth of Lévy trees. *Probability Theory and Related Fields*, 139(3-4):313–371, November 2007.

[Eva08] Steven N. Evans. *Probability and real trees*, volume 1920 of *Lecture Notes in Mathematics*. Springer, Berlin, 2008.

[FS09] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, Cambridge, 2009.

[GPW09] Andreas Greven, Peter Pfaffelhuber, and Anita Winter. Convergence in distribution of random metric measure spaces (Λ-coalescent measure trees). *Probability Theory and Related Fields*, 145(1-2):285–322, September 2009.

[Gro07] Misha Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, english edition, 2007. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[HM] Olivier Hénard and Pascal Maillard. On quasi-stationary distributions of subcritical Galton–Watson processes. In preparation.

[HP11] Chris Haulk and Jim Pitman. A representation of exchangeable hierarchies by sampling from real trees. *arXiv:1101.5619 [math]*, January 2011.

[Jan11] Svante Janson. Poset limits and exchangeable random posets. *Combinatorica*, 31(5):529–563, November 2011.

[Rém] Jean-Luc Rémy. Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire.(french)[an iterative procedure for enumerating binary trees and its application to their random generation] rairo inform. théor. 19 (1985), no. 2, 179–195. *MR0803997 (87h: 68132)*.

[Sta97] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997.

[Ver85] Wim Vervaat. Sample Path Properties of Self-Similar Processes with Stationary Increments. *The Annals of Probability*, 13(1):1–27, February 1985.